TECHNISCHE UNIVERSITÄT BERLIN

Hands-On PostgreSQL Demonstrations for Teaching and Learning Database System Internals

Stefan Halfpap

PGConf.DE | May 9, 2025

Database System (DBS) Internals are Complex

Relational database systems implement fundamental concepts

- Relational model and query optimization
- Persistent storage and caching
- Transaction processing

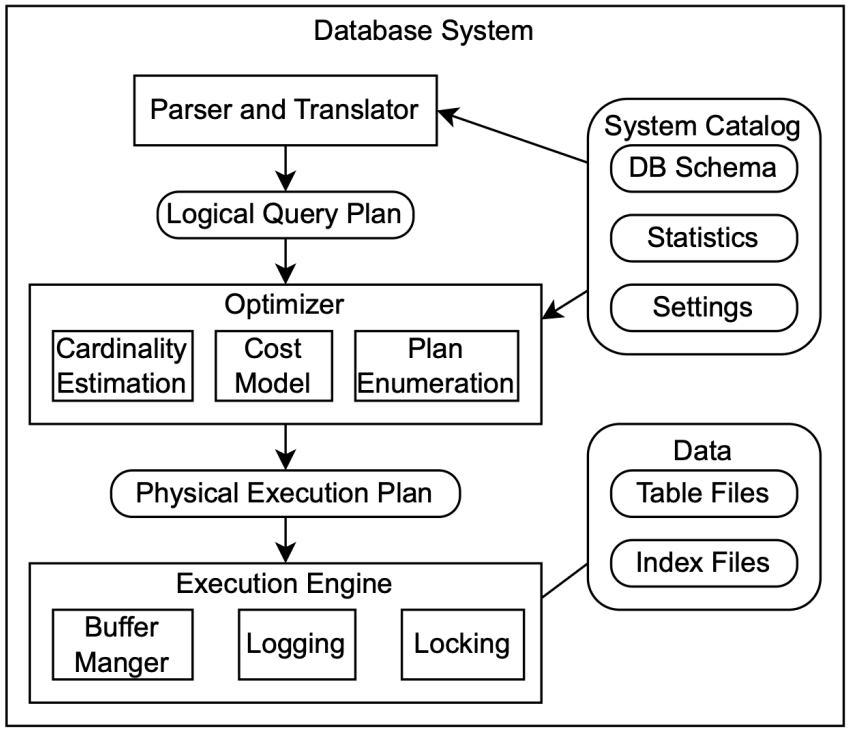

Simplified query processing

Good performance for diverse workloads

Learning DBS Internals is an Essential Skill

Database systems are omnipresent

Database systems include timeless patterns for performant system programming

→ If you can write code for a database system,
you can write code for almost any software system

How to Learn Database Systems Internals?

Database system courses and books cover concepts

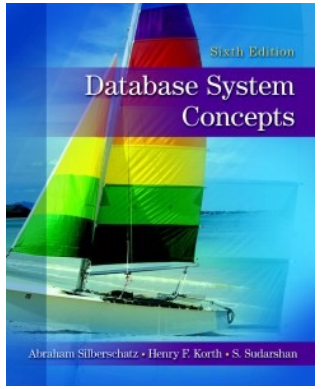

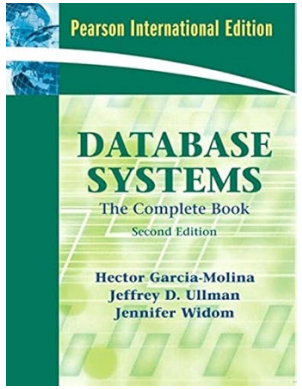

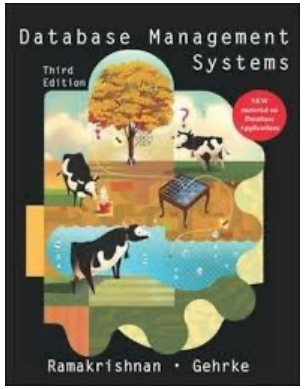

Learn theoretical concepts → ? → Learn practical system design and implementation

How to Learn Database Systems Internals?

Learn theoretical concepts ? → Learn practical system design and implementation

Read system documentation

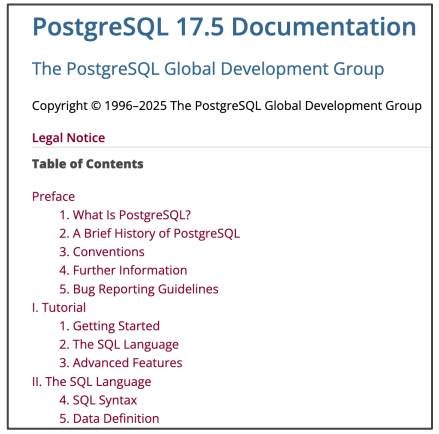
PostgreSQL 17.5 Documentation

The PostgreSQL Global Development Group

Copyright © 1996–2025 The PostgreSQL Global Development Group

Legal Notice

Table of Contents

Preface
- 1. What Is PostgreSQL?
- 2. A Brief History of PostgreSQL
- 3. Conventions
- 4. Further Information
- 5. Bug Reporting Guidelines

I. Tutorial
- 1. Getting Started
- 2. The SQL Language
- 3. Advanced Features

II. The SQL Language
- 4. SQL Syntax
- 5. Data Definition

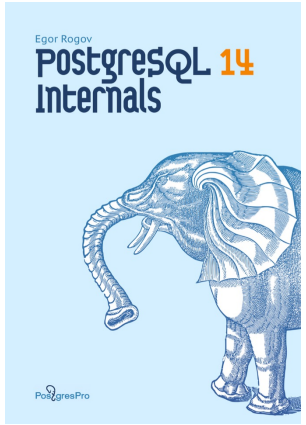

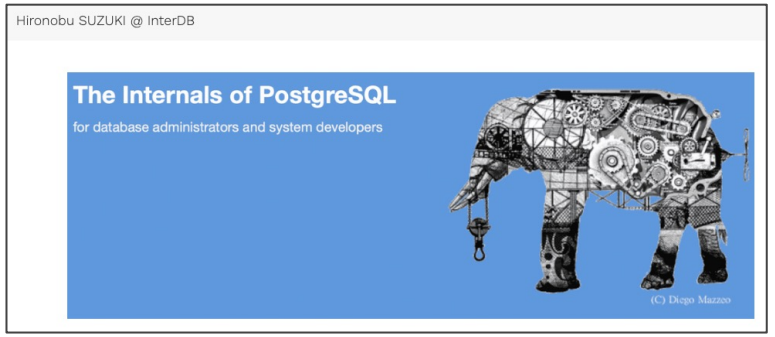

Demonstrations

Implementation-centric courses

Why PostgreSQL

- Popular, commonly-used, and mature system → relevant systems to know
- Great documentation
  (including books, blog posts, …)

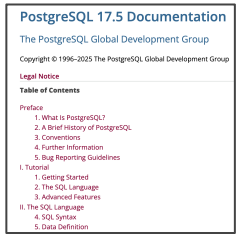

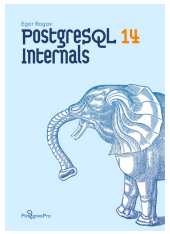

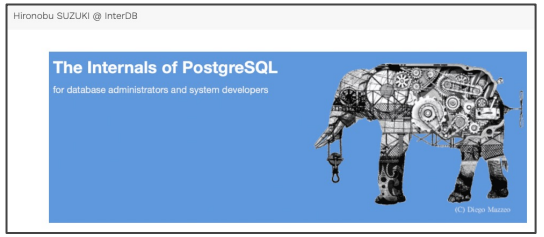

- Offers extensive capabilities for inspecting internals
  e.g., system catalogs, pageinspect, pg_buffercache, pgrowlocks, pg_walinspect
- Open-source → entry to deeper analysis

Collection of PostgreSQL Hands-on Demonstrations

The demonstrations are inspired, enabled, and partly adopted from existing write-ups and PostgreSQL extensions

Broad range of fundamental topics

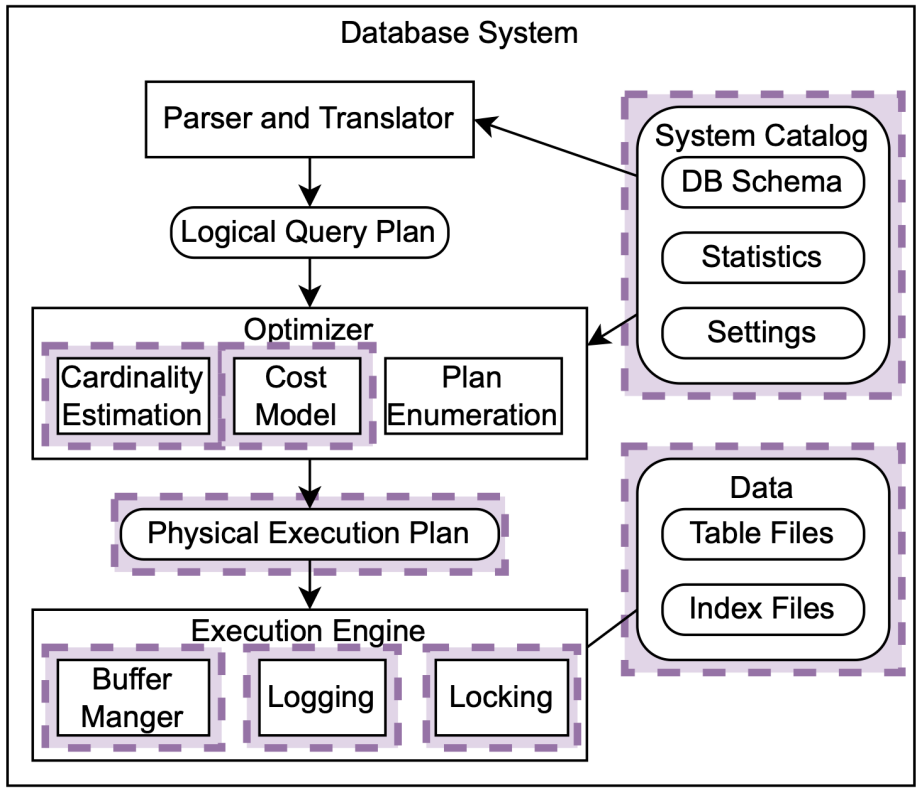

Lecture Topic	Suitable Demonstrations
Database Representation	File organization (Slotted) Page structure Tuple representation
Caching	Buffer inspection
Indexing	Index representation on disk Index utilization (sequential scan vs. index scan vs. bitmap scan)
Query Execution	Physical query plan Buffer usage
Query Optimization	Statistics Cardinality estimation Cost model Query plan inspection
Concurrency Control	Inspect running transactions Inspect MVCC columns Row locking
Recovery	Inspect running transactions Inspect the write-ahead log

https://github.com/klauck/demo_dbs_internals

How to use it: see https://github.com/klauck/demo_dbs_internals

1. Get the Scripts and Data

```
git clone https://github.com/klauck/demo_dbs_internals.git
cd demo_dbs_internals
```

2. Create and Start the Container

Ensure you **execute the following command from the root folder of the Git repository** to provide access to scripts and data:

```
docker run --name demo_postgres \
-v .:/root \
-e POSTGRES_USER=postgres \
-e POSTGRES_HOST_AUTH_METHOD=trust \
-e POSTGRES_DB=demo_db_internals \
-p 5432:5432 \
-d postgres:17
```

3. Load TPC-H Data

We provide a script to load TPC-H data with a scale factor of 0.01:

```
sh scripts/load_tpch.sh
```

How to use it: see https://github.com/klauck/demo_dbs_internals

4. Connect to PostgreSQL To connect using `psql`:

```
docker exec -it demo_postgres psql -U postgres -d demo_db_internals
```

⋮

Attribute-level statistics can be queried using the view `pg_stats`:

```
SELECT null_frac, n_distinct, most_common_vals, most_common_freqs, correlation
FROM pg_stats
WHERE tablename = 'nation' and attname = 'n_regionkey';
```

```
 null_frac | n_distinct | most_common_vals |   most_common_freqs    | correlation
-----------+------------+------------------+------------------------+-------------
         0 |       -0.2 | {0,1,2,3,4}      | {0.2,0.2,0.2,0.2,0.2} |   0.3476923
```

Hands-On PostgreSQL Demonstrations for Teaching and Learning Database System Internals

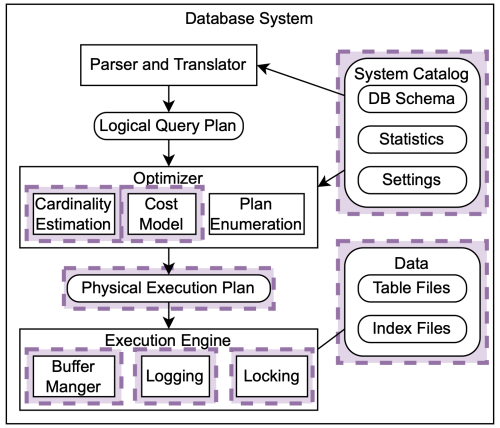

Lecture Topic	Suitable Demonstrations
Database Representation	File organization (Slotted) Page structure Tuple representation
Caching	Buffer inspection
Indexing	Index representation on disk Index utilization (sequential scan vs. index scan vs. bitmap scan)
Query Execution	Physical query plan Buffer usage
Query Optimization	Statistics Cardinality estimation Cost model Query plan inspection
Concurrency Control	Inspect running transactions Inspect MVCC columns Row locking
Recovery	Inspect running transactions Inspect the write-ahead log

Setup

You can run the demonstrations on an existing PostgreSQL installation. However, some experiments may require **superuser** privileges. Also, the **data-loading scripts must be modified** if you are not using Docker.

Docker Setup

Alternatively, you can run the demonstrations using Docker.

1. Get the Scripts and Data

```
git clone https://github.com/klauck/demo_dbs_internals.git
cd demo_dbs_internals
```

2. Create and Start the Container

Ensure you **execute the following command from the root folder of the Git repository** to provide access to scripts and data:

```
docker run --name demo_postgres \
-v .:/root \
-e POSTGRES_USER=postgres \
-e POSTGRES_HOST_AUTH_METHOD=trust \
-e POSTGRES_DB=demo_db_internals \
-p 5432:5432 \
-d postgres:17
```

3. Load TPC-H Data

We provide a script to load TPC-H data with a scale factor of 0.01:

```
sh scripts/load_tpch.sh
```

Attribute-level statistics can be queried using the view `pg_stats`:

```
SELECT null_frac, n_distinct, most_common_vals, most_common_freqs, correlation
FROM pg_stats
WHERE tablename = 'nation' and attname = 'n_regionkey';
```

```
 null_frac | n_distinct | most_common_vals |   most_common_freqs    | correlation
-----------+------------+------------------+------------------------+-------------
         0 |       -0.2 | {0,1,2,3,4}      | {0.2,0.2,0.2,0.2,0.2} |   0.3476923
```

https://github.com/klauck/demo_dbs_internals

Next Steps

- Get feedback
- Extend (Notebook; details for cost model and plan enumeration; visualization)

Stefan Halfpap
halfpap@tu-berlin.de

Image Attribution

Book Covers:

Silberschatz, A., Korth, H. F., & Sudarshan, S. (2010). Database System Concepts (6th ed.). McGraw-Hill Education.

Garcia-Molina, H., Ullman, J. D., & Widom, J. (2008). *Database Systems: The Complete Book* (2nd ed.). Upper Saddle River, NJ: Pearson Education.

Ramakrishnan, R., & Gehrke, J. (2002). *Database Management Systems* (3rd ed.). McGraw-Hill Education.

Rogov, E. (2023). *PostgreSQL 14 Internals*. Moscow: Postgres Professional. ISBN 978-5-6045970-4-0.

Screenshots:

The PostgreSQL Global Development Group: PostgreSQL 17 Documentation"
Retrieved from https://www.postgresql.org/docs/17/index.html on May 9, 2025.

Hironobu Suzuki: The Internals of PostgreSQL
Retrieved from https://www.interdb.jp/pg/ on May 9, 2025.

All images used under fair use for educational purposes.